Data Communications Requirements for CAS Applications

James D. McCabe

NAS Systems Development

jmccabe@nas.nasa.gov

Abstract

This paper discusses the long-haul data communications requirements needed to support visualization of Computational Aeroscience (CAS) simulations. A scientific computing model is presented, from which current communications requirements are derived. This model is then expanded to describe the CAS Collaborative Visualization Environment, in which scientists from multiple NASA centers will collaborate on multidisciplinary, time-accurate simulations. An analysis of this model, along with data from the TFCC Workload Model, shows that CAS data communications requirements will increase to 10-100 megabytes/second by 1996.*

1.0 Introduction

Visualization of Computational Fluid Dynamics (CFD) data is an increasingly important tool in the analysis of fluid flows for aeroscience simulations. As CFD simulations become more complex, such as in multidisciplinary or time-accurate simulations, visualization is necessary to help comprehend the large data sets generated.

*. TeraFLOPS Computing Configuration Model of System Parameters.

NASA's Computational Aerosciences (CAS) Project of the High Performance Computing and Communications (HPCC) Program is at the forefront in developing multidisciplinary, time-accurate simulations of fluid flows. In support of this project, applications for distance-independent visualization of large data sets are being developed[1].

A distributed client-server model for CAS visualization applications is emerging[2]. Like other client-server models, the visualization model depends greatly on a network infrastructure to provide data communications between components of the model. As distances between components increase, so does the importance of the data communications network.

In order for long-haul data communications networks to provide adequate support for CAS visualization applications, the communications requirements of these applications need to be determined. In addition, partitioning the visualization model to optimize long-haul communications between components is desirable.

Section 2 gives a background on the visualization of CFD simulations, including: an overview of the CAS Project and its impact on data communications support; a discription of the visualization process as a series of descrete components; partitioning of the computing/visualization system and analysis of the various partitions; and development of a model to describe communications requirements. Section 3 describes the current visualization model and data communications requirements derived from the model. Section 4 expands this visualization model to include the future CAS Collaborative Visualization Environment, and develops requirements to support that environment.

2.0 Background

There are three stages in CFD research: grid generation, numerical simulation of fluid flow, and post-processing of flow solution data[3]. A grid is generated to describe an object and its surrounding space. Flow solvers calculate the physical properties of the flow within the grid space, generating either a single solution file for steady-state flow or a solution file for each time step in a time-accurate flow*. Solution files contain values for energy, density, and momentum at grid points in the flow. Other physical properties may be derived from these values.

Post-processing of flow solution data consists of extracting the data of interest from the solution, applying the visualization technique that provides the desired interpretation of the data (i.e., isosurfaces, streaklines, or particle traces/rakes), generating geometries for the visualization technique, rendering and displaying the resulting view. For time-accurate solutions, animation of the views is desirable. CAS CFD simulations will be made more complex by their multidisciplinary nature.

2.1 Computational Aerosciences Project

Three computational problems have been chosen for the CAS Project: the High-Speed Civil Transport (HSCT), the High-Performance Aircraft (HPA), and the Advanced Subsonic Civil Transport (ASCT)[4]. Simulations of these problems will be multidisciplinary as well as time-accurate, composed of subsets from several scientific disciplines including:

- Computational Fluid Dynamics (CFD) modeling of airflow over the aerospace vehicle.
- Structural analysis of the vehicle.
- Analysis of the control surfaces of the vehicle.
- Thermodynamics.
- Acoustics.
- Chemistry.

CAS Scientists from NASA research centers (Langley, Lewis, and Ames Research Centers) will collaborate to visualize the results of these multidisciplinary, time-accurate simulations. New distributed visualization applications, such as the *distributed Virtual Wind Tunnel* or *Tempus Fugit/Interview*, are being developed in part to support this collaborative visualization environment.[5,6]

It is expected that a distributed visualization model will be used between the CAS centers to support this collaborative effort. This model will take a client-server approach to visualization, central-

---

*. Physical parameters do not change with time for steady-state flows, but do change with time for time-accurate (or unsteady-state) flows.

izing components that are compute or memory-intensive, while distributing components that interact with the user. Since distances between the CAS centers are cross-country (Fig. 1), the ability of long-haul networks to meet the data communications requirements between components of the distributed visualization model is crucial.

FIGURE 1. NASA Centers Working on CAS Computational Problems

2.2 Visualization Process

Visualization is the process of transforming data into a format that facilitates comprehension on a macroscopic level. For CAS problems there are large amounts of data (currently up to hundreds of gigabytes) requiring media such as pictures (still or animated) to allow a human to absorb the data and determine what, if any, sections require closer scrutiny. By iterating this process, the scientist can focus on the most interesting parts of the problem. The visualization process can be separated into several components which act in a serial fashion.

Solution data for the discipline(s) of interest are generated using a flow solver, typically on a supercomputer or parallel system, and are based on input data from the scientist in the form of grids and initial values. For time-accurate simulations, solution data is generated for each time step. Current time-accurate simulations range from hundreds to thousands of time steps. This solution data, in the form of *solution files*, is then transferred to and stored on disk.

Subsets of a solution's scalar and vector fields may be selected from the solution file(s). These subsets are based on *extracts*[7]. A geometry processor then receives the solution file(s) or subset(s), called here *post-processing files*, and creates the geometries needed for display of the object* and flow field characteristics, based on the visualization technique chosen by the scientist. This *geometry data* is transferred to a rendering processor which, with input about the desired view (viewing distance/angles) from the scientist, processes (polygon scanning, shading) the data and produces an image suitable for display. At the end of the visualization process monitors or virtual reality devices are used to view the solution.

This process may be modeled as shown in Fig. 2.

FIGURE 2. Visualization Model

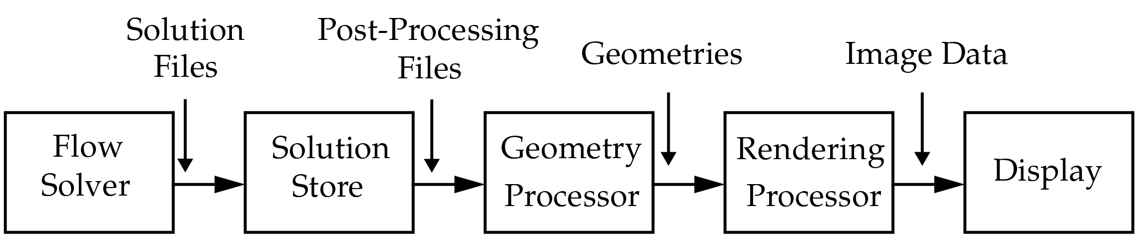

This model is based on VanZandt's model of scientific computing.[8] When this model is coupled with computing/visualization systems and communications networks, partitions within the model can be mapped to communications requirements.

2.3 Partitioning

Gerald-Yamasaki illustrated how various partitions in the computing/visualization system determined the ability of the system to provide an effective environment for cooperative visualization.[9]

*. For CAS simulations, the object is typically an aerospace vehicle.

Consider a basic computing/visualization system, such as a supercomputer and graphics workstation, connected via a network. Figure 3 shows how the visualization model may be partitioned across this system:

FIGURE 3. Computing/Visualization System

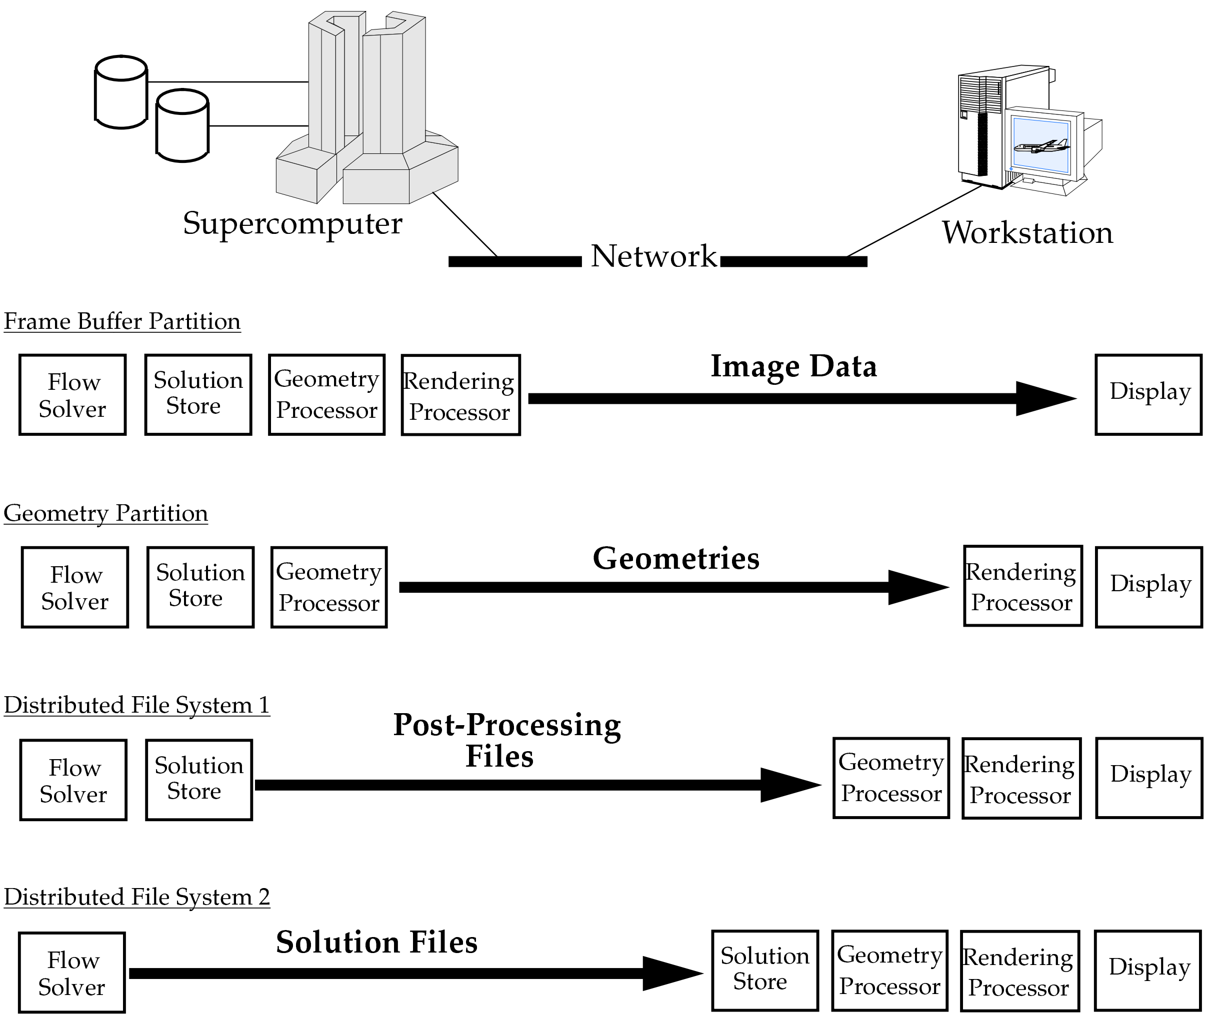

In each of these partitions the network is used to connect components of the visualization model through communications that have different capacity and delay characteristics than those encountered in other parts of the computing/visualization system (computer backplane, disk and display I/O channels). For a high-performance local-area network, connecting systems in a computer room environment, capacity and delay characteristics approach or surpass those for other components of the system. When the network extends to the metropolitan or wide-area environments, however, capacity and delay characteristics are currently inferior to other components of the system. It is therefore important to understand the effectiveness of each partition in the long-haul communications environment.

In this section, general characteristics of each partition are given. Detailed analyses of communications requirements are given in sections 3 and 4.

Frame Buffer Partition

In this partition the supercomputer performs all of the computing and visualization functions, and the workstation (or frame buffer) acts solely as a display device. Image data in the form of pixels is transferred across the network. Network capacity requirements are determined by display and pixel information sizes.

Geometry Partition

The geometry partition focuses on two types of data distribution: transferring geometry data generated on the supercomputer to the workstation for rendering and display; or distributing the geometry processing between the supercomputer and workstation via a distributed graphics library. Network capacity requirements are based on the visualization technique used, size and sophistication of object viewed, and wheter flow is steady-state or time-accurate.

Distributed File System Partitions

These partitions are based on the transfer of whole or partial solution files into or out of mass storage. In the case where data is transferred from storage to the geometry processor, as in *Distributed File System 1*, network capacity is required to sustain the processing capability of the geometry processor. For simulations of time-accurate flows, an input rate to the geometry processor of 10 post-processing files/second is expected. In *Distributed File System 2* whole solution files are transferred across the network to storage. Capacity requirements are based on file size and transfer time. For solutions of time-accurate flows, compute times can be used as a first-level approximation of transfer times. As the solution for timestep *N+1* is being generated on the supercomputer, the solution for timestep *N* can be transferred across the network.

These partitions of the computing/visualization system, while simplistic, are a start towards taking a systems-wide approach to determining communications requirements.

2.4 Communications Requirements Model

Capacity and delay characteristics can be combined to define a communications performance envelope whereby the communications requirements of various applications can be compared relative to each other and to performance levels of communications technologies/services. Figure 4 shows the generic requirements envelope that will be used for CAS visualization applications.

FIGURE 4. Communications Requirements Envelope

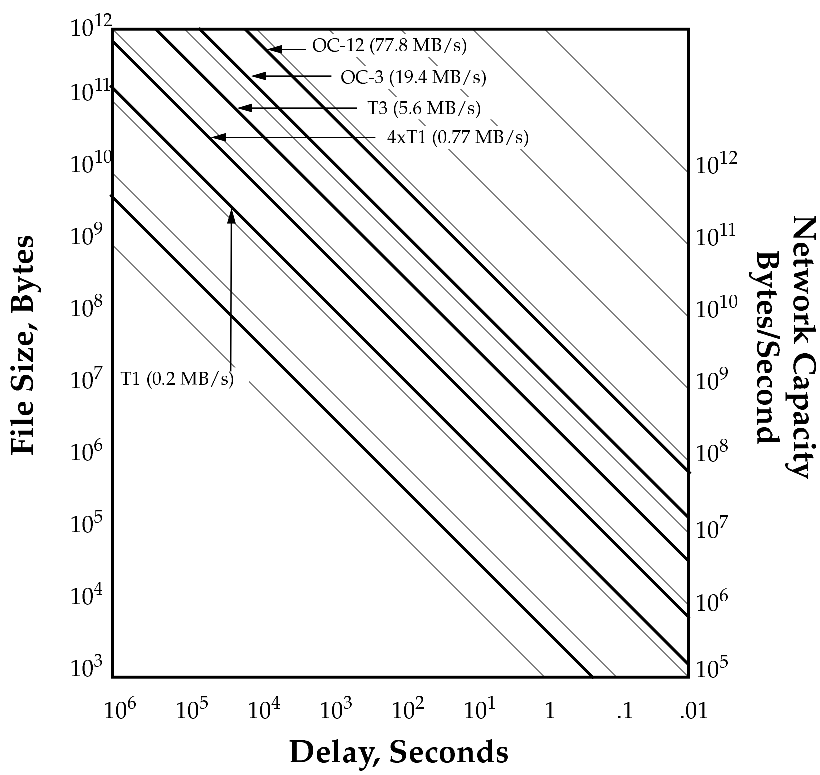

In this figure, diagonal lines represent network capacities. Several T-carrier and Synchronous Optical NETwork (SONET) service levels are shown as solid diagonal lines. These services fall within the capacity range expected for CAS requirements.

3.0 Current Data Communications Requirements

Current requirements are derived from the VanZandt model shown in Section 2.2, using the following assertions:

- A single scientific user, single- or multiple-system model is applied.
- Visualizations are of single discipline, steady-state simulations.

Partitions of this model were described in Section 2.3. An analysis of these partitions is done with data from the NAS TFCC Workload Model.

3.1 Analysis

Frame Buffer Partition

Capacity requirements for the frame buffer are determined by display and pixel information sizes. Typical displays currently use a frame size of 1024 by 1280 pixels which, with 3 bytes of data per pixel, is 3.9 MB of data per frame. At an average display rate between 10 and 30 frames/second, depending on the level of interaction with the frame buffer and method of data acknowledgement used, delay characteristics range from 0.1 to 0.03 seconds. This is neglecting any protocol overhead.

Geometry Partition

For this model the geometry partition is based on transferring geometry data generated on the computation system (supercomputer) to a single visualization system (graphics workstation). Applications that follow this model use a mechanism such as *Remote Procedure Call* (RPC) to transfer data, as necessary, to the visualization system. Current geometry data sets range from 10^4 to 10^7 Bytes in size and have delay characteristics from 0.1 to 1 second. Note that the geometry data sets transferred are subsets of the entire geometry data set, and are transferred as required by the visualization system.

Distributed File System Partitions

Visualization applications currently use the *Distributed File System 2* partition for transfer of solution files from the supercomputer to graphics workstation. Since solutions considered here are steady-state, the computation delay discussed in Section 2.3 is not a factor. Current solution file sizes range from 10^7 to 10^{10} Bytes, with file transfer times of up to 10^4 seconds. The *Distributed File System 1* partition is not currently used.

These requirements are now be mapped onto the communications model presented in Section 2.4.

FIGURE 5. Current Requirements Mapped Onto the Communications Model

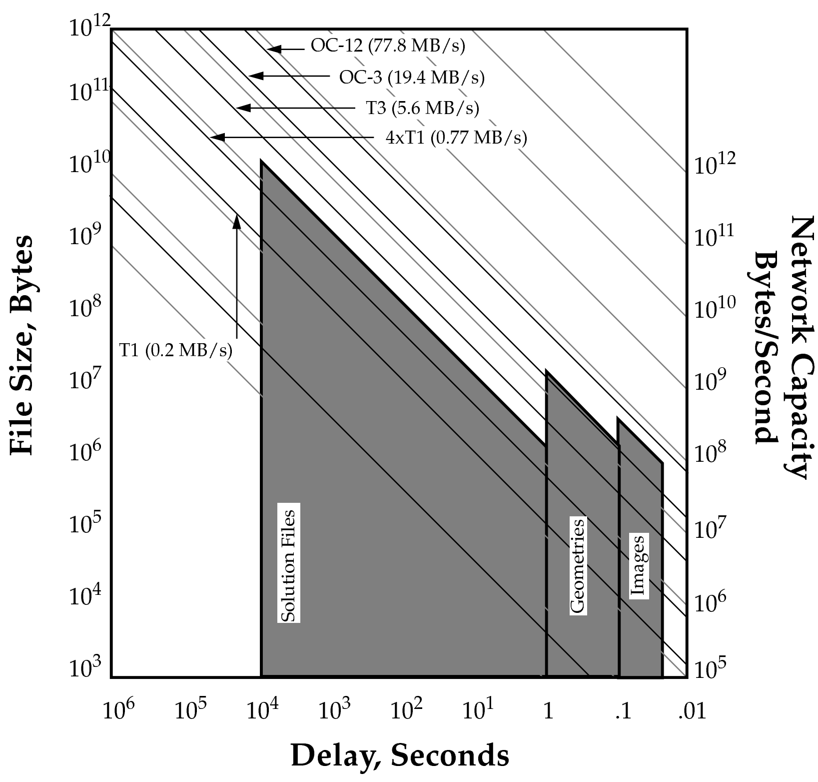

T-carrier services (up to T3) have sufficient capacity for current solution file transfer requirements and most geometry and image requirements, but the upper limits of geometry and image transfer requirements are in the SONET OC-3 to OC-12 range, which are outside current capabilities. Solution file and geometry transfer delay and size characteristics map well to current long-haul capacities, whereas image transfer characteristics are not well suited to long-haul transmission.

4.0 Future Data Communications Requirements

VanZandt's scientific computing model is now modified to represent the collaborative visualization environment needed for multidisciplinary, distance-independent visualizations. CAS data communications requirements for this environment are based on:

- Multiple scientific users distributed among the NASA Centers shown in Section 2.1.
- A distributed client-server model for CAS visualization applications (visualization server).
- Visualizations of multidisciplinary, time-accurate simulations.

These assertions form the foundation for the future CAS Collaborative Visualization Environment.

4.1 CAS Collaborative Visualization Environment

It is expected that the scientists should have the ability to collaborate at various levels of the simulation. For example, scientists may want to simultaneously view and manipulate the same visualization, run multiple independent visualizations of the same simulation, or share modifications to the simulation. In addition, scientists will want audio and visual communications with each other. A visualization server model, distributed across the NASA Centers, shows promise towards supporting such collaborative efforts. Emerging communications technologies and services are expected to support this model, as well as provide audio and video communications. Figure 6 shows how the CAS Collaborative Visualization Environment may be deployed.

FIGURE 6. CAS Collaborative Visualization Environment

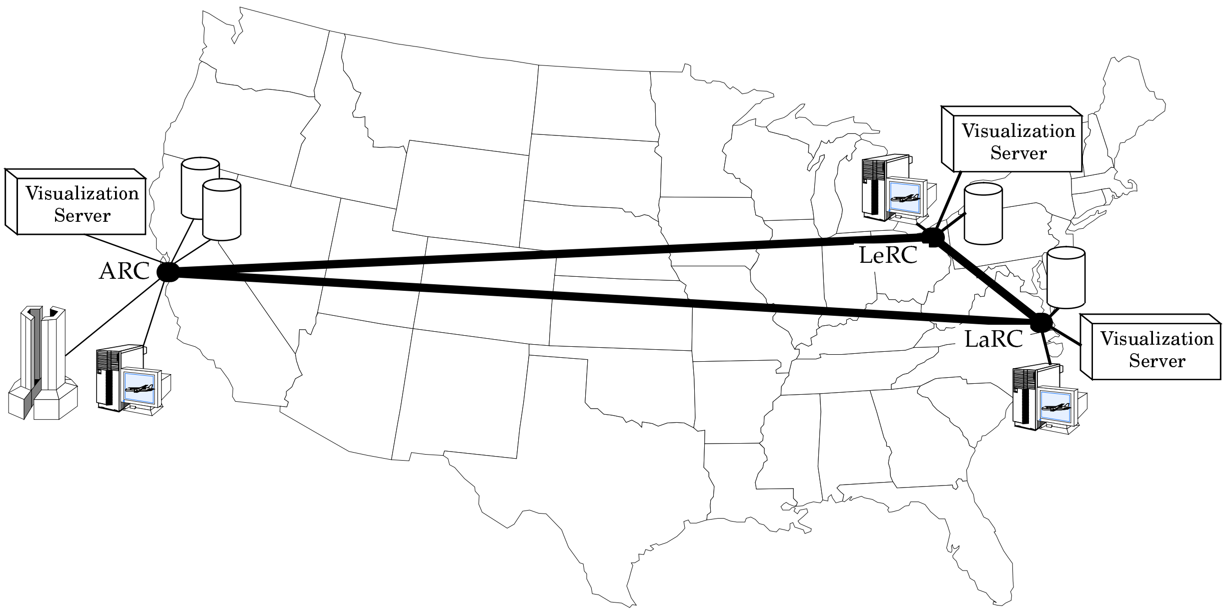

In this figure, visualization servers and mass storage systems are deployed at all centers, while computing systems (supercomputers and massively parallel systems) are localized to a single center. This follows from the trend of supercomputing consolidation that NASA and other groups in government and industry are following. Greater amounts of storage will need to be co-located with the computing systems, as a solution file cache or staging area for distribution to the other centers.

With the deployment of multiple storage and visualization systems at the centers comes an increase in the complexity of communications between the systems. Along with the compute system-workstation data transfer model which forms the basis for current CAS communications requirements, data transfer to and from mass storage and visualization servers also needs to be considered. Simultaneous data communications to, from, and between multiple workstations is also a part of this environment.

4.2 CAS Collaborative Visualization Model

Figure 7 illustrates how the VanZandt model is expanded to include the CAS Collaborative Visualization Environment. Important features of the expanded model are: communications between each step *N* and all elements of step *N+1* (for all solution store, geometry and rendering processors); communications between all elements within each of the geometry and rendering processor steps; and staging of the solution files from the flow solver to local storage, then from local storage to remote storage*. It is important to note that this expanded model does not cover all possible communications flows. Only the most likely flows are shown.

FIGURE 7. CAS Collaborative Visualization Model

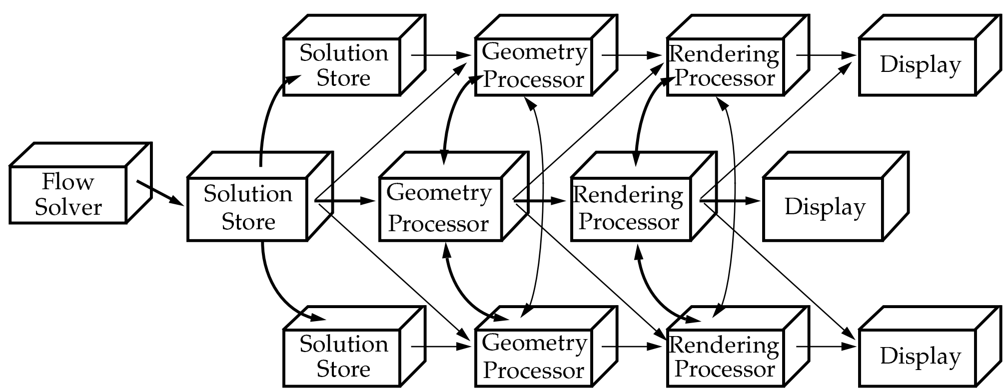

Partitions for this model, along with the CAS collaborative computing/visualization system, are given in Figure 8. In this figure, horizontal arrows represent flows from step *N* to step *N+1*, while curved arrows represent flows within a step. In *Distributed File System 2*, horizontal arrows between solution stores represent staging of solution files from the local to remote systems.

*. “Local” and “Remote” refer to locations relative to the compute systems.

FIGURE 8. CAS Collaborative Computing/Visualization System and Partitions

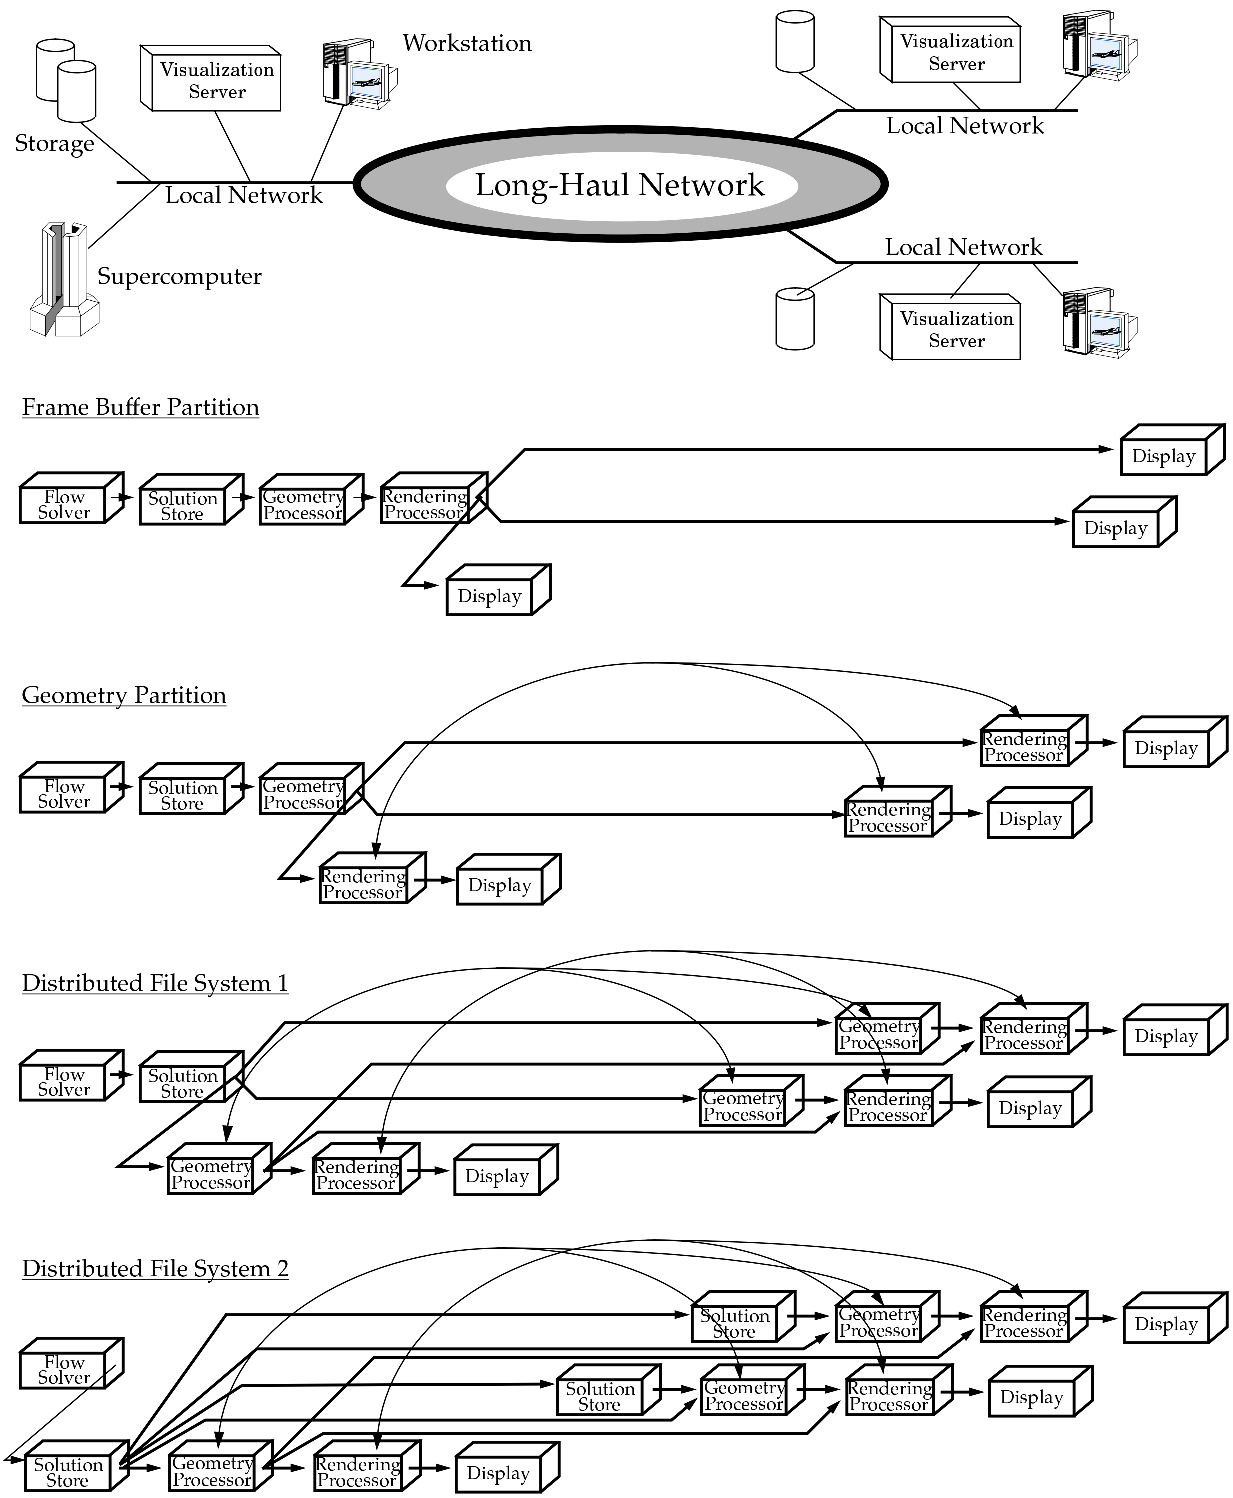

4.3 Analysis

Inputs are made for future communications requirements from the TFCC Workload Model for time-accurate CAS solutions, using the visualization model from Section 4.2. These inputs are averages of grid sizes and number of time steps saved (See Table 1).

TABLE 1. TFCC Workload Model Inputs to Visualization Model

	Avg Grid Points	Avg Timesteps
FY94	2890000	377
FY95	4000000	478
FY96	4050000	552

Frame Buffer Partition

From the analysis of Section 3.1, the frame buffer partition was found not to scale well across a long-haul network. This was due to the combination of large data size (3.9 MB) coupled with strict delay characteristics (0.1 to 0.03 seconds). Although this partition is not likely to be used in the CAS Collaobrative Visualization Environment, an analysis is included here for informational purposes. Future devices may include holographic displays. Frame size for a holographic display that is 512 pixels on an edge would be 512^3, or 134 Mpixels. At 3 Bytes of data per pixel, a frame would be 402 MB in size, which, at 30 frames/second, requires over 12 GigaBytes/second (GB/s) of network capacity. This is greater than the highest SONET data rate specification, and far beyond the scope of available network capacity in the timeframe of CAS application deployment.

Geometry Partition

Expansions to the geometry partition for the CAS Collaborative Visualization Environment include the transfer of geometry sets to multiple rendering processors. Two modes of operation with this partition are envisioned; a client-server relationship between workstations may exist, where one workstation maintains viewing control and other systems act as passive viewing devices; or each workstation acts in an independent, yet sychronized, fashion. Transfers may be via a multicast mechanism if synchronization between workstation is required. Rendering information may be passed between rendering processors to support synchronization for multiple simultaneous visualizations or whenever visualizations interact with each other.

Delay, file size, and capacity requirements for geometry data transfer are given in Table 2.

TABLE 2. Delay, File Size and Network Capacity Requirements for Geometry Data Transfer

	Geometry File Size (MBytes)	Delay (seconds)	Network Capacity (MB/sec)
FY94	306	1	306
FY95	482	1	482
FY96	561	1	561

This data is based on a single user application. For the CAS collaborative visualization environment, this data is multiplied by *N*, which represents the number of simultaneous users and is based on the mode of operation. If the data is multicast over the long-haul network, a multiplier may not be needed. In comparison to geometry data requirements, control information passed between rendering processors is expected to have relatively low capacity requirements (0.001 - 0.01 MB/sec).

Distributed File System Partitions

CAS visualization applications will use both *Distributed File System 1 and 2* partitions for data transfer. As discussed in Section 2.3, *Distributed File System 1* will be used when post-processing files are simultaneously fed to multiple geometry processors, to support multiple independent visualizations. *Distributed File System 2* will be used to stage solution files from local to remote storage.

TABLE 3. Requirements for *Distributed File System 1*

	Solution File Size (MBytes)	Timesteps	Post Processing File Size (MBytes)	Delay (seconds)	Network Capacity (MB/sec)
FY94	34865	377	92.5	0.1	925
FY95	61184	478	128	0.1	1280
FY96	71539	552	129.6	0.1	1296

TABLE 4. Requirements for *Distributed File System 2*

	Solution File Size (MBytes)	Delay (seconds)	Network Capacity (MB/sec)
FY94	34865	10^4	3.5
FY95	61184	10^4	6.1
FY96	71539	10^4	7.2

Delays for each partition are based on the requirments of the storage systems or geometry processors receiving the data. For *Distributed File System 1*, delay is based on an imput rate of 10 frames/second to the geometry processor, while for *Distributed File System 2*, delay is based on allowable time to stage files from local to remote storage. This time was chosen to be equivalent to current file transer times.

These requirements are used to modify the Communications Model presented in Section 3.1.

FIGURE 9. Future CAS Requirements Added to the Communications Model

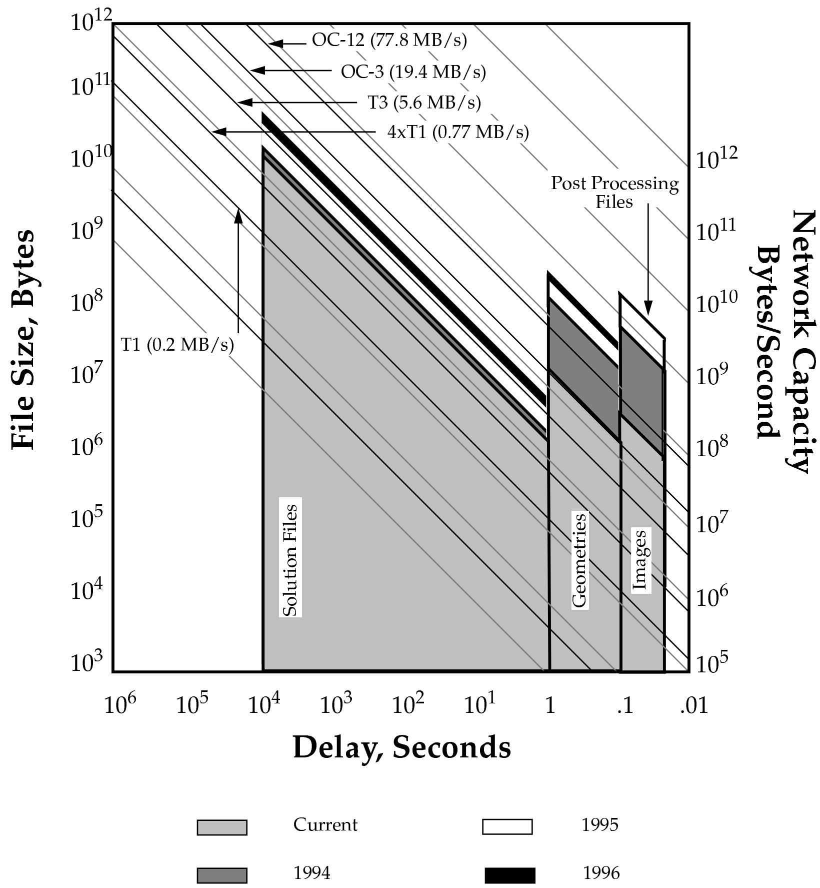

Requirements for *Distributed File System 1* (Post Processing Files) and *Frame Buffer* (Images) partitions are beyond the scope of planned long-haul network capacity/delay characteristics. Requirements for *Distributed File System 2* (Solution Files) and *Geometry* (Geometries) partitions will drive the deployment of new data communications technologies.

From this figure, by 1996 long-haul network capacity needs to increase to about 10 MBytes/second to support solution file transfers, and to greater than 100 MBytes/second to support geometry data transfers.

5.0 Conclusions

This paper discussed the long-haul data communications requirements needed to support visualization of Computational Aeroscience (CAS) simulations. Current communications requirements were derived from the VanZandt model of scientific computing, applied to a single user, supercomputer-workstation system visualizing a single discipline, steady-state simulation. Under these conditions, the model was partitioned to show how the long-haul network would fit into the system.

This model was then expanded to describe the CAS Collaborative Visualization Environment, in which scientists from multiple NASA centers will collaborate on multidisciplinary, time-accurate simulations. An analysis of this model, along with data from the TFCC Workload Model, showed that:

- Capacity and delay requirements of the *Frame Buffer* and *Distributed File System 1* partitions do not map well to the long-haul environment.
- Geometry and Distributed File System 2 partition requirements will drive the deployment of high-performance long-haul communications technologies for CAS work. In particular, by 1996 CAS long-haul data communications requirements will increase to 7 MBytes/second for solution file transfer and over 100 MBytes/second for geometry data transfer.

References

1. Bailey, F.R., and Simon, H.D., Future Directions in Computing and CFD. NAS Technical Report RNR-92-019, May 1992.
2. Rogers, S.E., Buning, P.G., and Merritt, F.J., Distributed Interactive Graphics Applications in Computational Fluid Dynamics, *International Journal of Supercomputer Applications*, Vol. 1, No. 4, Winter 1987, pp.96-105.
3. Simon, H.D., Van Dalsem, W., and Dagum, L., Parallel CFD: Current Status and Future Requirements. In *Parallel CFD - Implementations and Results Using Parallel Computers*. MIT Press, Cambridge, Mass., 1992.
4. Claus, R.W., Holst, T.L., Salas, M.D., The NASA Computational Aerosciences Program-Toward TERAflops Computing. *Proceedings of the 30th Aerospace Sciences Meeting & Exhibit*, January 1992.
5. Bryson, S., and Gerald-Yamasaki, M., The Distributed Virtual Windtunnel, NAS Technical Report RNR-92-010, March 1992.
6. Gerald-Yamasaki, M.J., Cooperative Visualization of Computational Fluid Dynamics. NAS Technical Report RNR-92-007, March 1992.
7. Globus, A., A Software Model for Visualization of Time Dependent Computational Fluid Dynamics Results., NAS Technical Report RNR-92-031, November, 1992.
8. Van Zandt, J., Scientific Computing Paradigm, NAS Technical Report RND-92-000, August, 1992.
9. Gerald-Yamasaki, M.J., Interactive and Cooperative Visualization of Unsteady Fluid Flows. NAS Technical Report RNR-92-018, March 1992.